

Research on Scientific Components at Argonne

Lois Curfman McInnes (PI), Jay Larson, Boyana Norris, and Jason Sarich
Mathematics and Computer Science Division
Argonne National Laboratory

Summary

The Center for Component Technology for Terascale Simulation Software (CCTTSS), a multi-institutional SciDAC-funded project, focuses on developing component technology for high-performance computing via the Common Component Architecture (CCA), with a goal of facilitating the interoperability and reuse of scientific software. Highlights of recent work by the Argonne contingent of the CCTTSS include developing scientific components, incorporating the CCA into climate simulations, improving the ease of use of CCA environments, and investigating issues in computational quality of service related to robust, efficient, and scalable performance.

Scientific Components. A key area of recent work among CCTTSS researchers at Argonne has been the development of production components that are used in scientific applications as well as prototype components that aid in teaching CCA concepts. Among the freely available components are tools for parallel “MxN” data redistribution, based on experience within the Model Coupling Toolkit, and tools for optimization and linear algebra. An important facet of this work is defining domain-specific common interfaces, which helps in realizing our vision of interchangeable scientific components. We are collaborating with the TOPS¹ SciDAC center to define common interfaces for linear and nonlinear solvers (see Figure 1).

Optimization in Quantum Chemistry. An important goal of our project is to assure that scientific components are interoperable, able to deliver high performance, and useful to

real applications. Thus, we are collaborating with chemists to perform CCA-based quantum chemistry simulations, which employ components based on the NWChem (PNNL) and MPQC (SNL) quantum chemistry codes for energy, gradient, and Hessian computations; the TaoSolver optimization component (ANL); and components based on Global Arrays (PNNL) and PETSc (ANL) for parallel linear algebra operations. Recent molecular geometry experiments have demonstrated reductions in simulation times up to 43% compared to the stand-alone chemistry packages.

Climate Components. Another Argonne application focus is computational climate modeling, which is critically important for our understanding of global processes and the potential for human impact. Recent work in climate applications has proceeded on multiple fronts. We continued to leverage SIDL/Babel to create a language-independent version of MCT, which will allow users to construct multilanguage

¹ Terascale Optimal PDE Simulations Center, PI: D. Keyes,
<http://tops-scidac.org>

parallel coupled models using a toolkit interface. We are in the process of creating a set of “parallel coupling” CCA components built by using MCT and CCA technologies for inclusion in a future release of the CCA toolkit. These parallel coupling components include MCT-based implementations of the new CCA common interface specifications for distributed arrays and parallel data redistribution, as well as MCT’s parallel data transformation utilities. We are continuing work on additional contributions to the CCA Toolkit, including formal release of a netCDF component and development of a CCA component interface to PnetCDF.

Computational Quality of Service.

Component-based environments provide opportunities to improve the performance and numerical accuracy of scientific simulations. The concept of the automatic selection and configuration of components to suit a particular computational purpose is called *computational quality of service (CQoS)*. Recent CQoS work at Argonne has focused on developing infrastructure to consistently collect and access both runtime and historical performance data. We implemented a prototype component infrastructure that supports performance monitoring, analysis, and adaptation of important numerical kernels, such as nonlinear and linear system solvers. We defined a simple, flexible interface for the implementation of adaptive nonlinear and linear solver heuristics. We also provide components for monitoring (based on TAU), checkpointing, and gathering of performance data.

Usability. Argonne has taken a leading role in a new effort focusing on the usability of high-performance component technologies. A highlight of recent work is the development of a new Eclipse-based integrated development environment (IDE) for CCA components (see Figure 1). In addition to the standard single-language IDE

support for configurable editors, multiple code views, build and execution support, the CCA IDE supports the development of language-independent interfaces, from which code in multiple languages, such as C, C++, Fortran, or Java can be generated. Multilanguage editing, browsing, and build support are some of the features that are not provided by traditional IDEs but are characteristic of CCA applications and are thus included in the CCA IDE implementation. In addition, capabilities provided by other Eclipse projects targeting high-performance computing, such as the parallel launch and debugging support of the Parallel Tools Platform (PTP), will be integrated seamlessly with the CCA IDE, providing support for the complete lifecycle of applications.

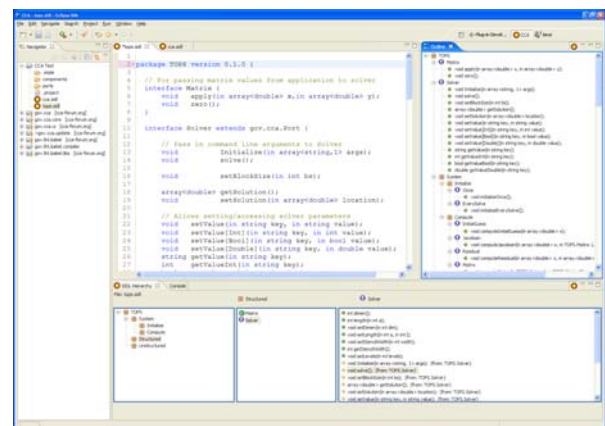


Figure 1 Screenshot of the Eclipse component development environment showing the TOPS nonlinear and linear solver component interfaces.

For further information on this subject contact:

Lois Curfman McInnes (PI)
Argonne National Laboratory
Argonne, IL 60439-4844
Phone: 630-252-5170
E-mail: mcinnes@mcs.anl.gov
<http://www.mcs.anl.gov/cca>